Large Participation Experiential Learning Activity for Multiagent Systems

Alan Tsang

Carleton University, Ottawa, Canada
alan.tsang@carleton.ca

Abstract

Multiagent systems is a key area within artificial intelligence (AI) that explores the behavior of interacting rational agents where the decisions of one agent impact others. Rooted in economic game theory, multiagent systems takes the idea of individual incentives from economic game theory and applies it to distributed computation and decentralized mechanisms. It examines not only how certain overall economic or computational goals can be accomplished, but also why individual participants will choose to cooperate with reaching that goal. While multiagent systems is grounded in rigorous mathematical theory, current pedagogical approaches often lack opportunities for students to connect abstract theory with real-world human dynamics. This disconnect is particularly pressing as AI increasingly operates in sociotechnical environments, where understanding human behavior and interaction is critical. This paper presents the first exploration of using large participation activities to facilitate experiential learning to bridge this gap. We report on a day-long resource allocation scenario involving up to 43 participants, designed to simulate multiagent interactions under pressure and with meaningful stakes, where learners can apply their theoretical knowledge to analyze and solve emerging problems. We propose "megagames" as a powerful pedagogical tool not only for multiagent systems, but also for other domains as well.

Introduction

Multiagent systems is a vibrant discipline within the wider umbrella of artificial intelligence (AI) that examines the behavior within communities of reasoning agents. These agents can be humans, organizations, or even algorithms. They reason in the sense that they make decisions in order to achieve well defined objectives (often to maximize some kind of utility function). These utility functions can encode a wide variety of behaviors, from selfish self-interest or hedonic utility, to altruism and empathy, to spitefulness. Agents can also be constrained by computational power, or be stochastic or error-prone.

Pedagogy within multiagent system is rich in mathematical rigor and theory. Small learning demos in the classroom can showcase simple game theory models, but these activities are more demonstrative in nature rather than experiential due to the lack of meaningful stakes for the participants. We identify a gap in pedagogical practices that ties training in theory with the application of that theory to actual human agents. This is especially important as AI is deployed in increasingly sociotechnical domains.

We present the first exploration of using large participation activities to facilitate experiential learning to bridge this gap in theory pedagogy. We report on an activity conducted in Fall 2022 and 2023, involving members of the community in both organization and participation. It is part of a standalone course on multiagent systems for computer science graduate students, and is crosslisted as an upper-year computer science undergraduate students. The course enrolls approximately 30 students each year, with around 5 graduate students. The course is structured with 9 weeks of lectures, followed by the experiential learning activity and then project presentations.

Background

Multiagent Systems

While artificial intelligence is generally concerned with building reasoning systems that solve well defined problems, multiagent systems examines the behavior of populations of such reasoning systems. One cornerstone of multiagent system is economic game theory which describes how patterns of behaviors can emerge using simple mathematical models. The most classic example in game theory is the Prisoner's Dilemma, a normal form game that imagines the players as two prisoners being interrogated by separately by the police. Each player can either Cooperate with the other (i.e. remaining silent), or Defecting (i.e. testifies). The best outcome a player can hope for is that she testifies (Defects) while the other player remains silent (Cooperates). In fact, regardless of what the other player does, each player would be better off testifying (i.e. it is a dominant strategy). The dilemma arises because if both players testify, the police have enough evidence to send them both in jail.

This dilemma is partially resolved when considering a scenario where the Prisoner's Dilemma is not a one-shot game, but imagines the two engaging in repeated play. Cooperative behavior can be sustained by the threat of immediate retaliatory action (a strategy called "Tit for Tat" or "TFT"). This game is often used as an example of how pure self interest can lead to cooperative behavior in a population.

A key principle in game theory is the idea that individuals behave according to incentives. Multiagent systems applies this idea of individual incentives to classical AI studies (Rosenschein 2025). Whereas classical AI systems are concerned with the "what" and "how" of computation, multiagent systems is concerned with the motivation of the individual components, which are described as agents (often human individuals or organizations, but recently also computational agents acting on behalf of them).

The study of auctions offers an example of how agent motivations are taken into account. In an auction, the seller wishes to allocate a single good (the sale item) to the person who values it the most (equivalently: to the person willing to pay the highest for it). The procedure of the auction can be constructed so that it is in the buyer's best interest to bid in a way such that their valuation for the item is revealed and the item sold to the buyer with the highest valuation. We can further expand the study of auctions from these relatively simple single-item auction to allocation of a complex variety of goods, including mechanisms where money is not involved and allocation fairness is the main desiderata (for example, for the national kidney exchange for UNOS (Abraham, Blum, and Sandholm 2007)). In the other direction, multiagent system is also concerned with the fair selection of a single or a few goods, in the form of the winner of elections. Recent exploration in this field include studies of strategic voting (Meir 2018), strategic timing (Tsang and Larson 2017), and partially decentralized mechanisms such as liquid democracy (Kahng, Mackenzie, and Procaccia 2021).

Experiential Learning

Experiential learning is a pedagogical paradigm (Kolb 2014) where learning is accomplished by direct observation and interaction, rather than by reading or traditional lectures. Here, learners must actively participate in the problem solving experience and then reflect on that experience using analytical skills to draw abstractions from the experience to be applied in future problems. The salience of the experiences make them memorable and lead to more durable learning. Moreover, the hands-on approach makes the relevance of the lesson readily apparent to the learner, and this motivation is a key driver of successful learning outcomes (Reigeluth 2013; Tohidi and Jabbari 2012; Lepper 1988).

One of the most successful experiential learning programs in multiagent systems is the RoboCup robotics competition where teams program robots that compete in football (Baltes, Sklar, and Anderson 2004; Vidal and Buhler 2002). Elsewhere in computer science, experiential learning has also been deployed to teach energy grid interdependencies (Scarlatos, Tomkiewicz, and Courtney 2013), operation of quadcopters (Lee et al. 2024), and awareness of predatory website design (Ye et al. 2025).

Some experiential learning approaches embed the experience within a competition or gamified framework; these are sometimes distinguished as Goal Based Learning or Serious Games. By framing the learning experience as a game, the learning task become fun and engaging. The learner is driven by innate curiosity or competitiveness. This intrinsic motivation can be more enduring than external drivers, and can be harnessed to improve the robustness of the learning experience (Deci and Ryan 2013; Lee and Hammer 2011). These framework have seen use in as a diverse domain from civil engineering (Ebner and Holzinger 2007), to learning Bayesian games (Burguillo 2010), and coding for AI robustness (Geleta et al. 2023).

Megagames

Megagames are sometimes described as a combination of Model UN (United Nations n.d.) and a tabletop board game[1]. They are large group activities involving dozens of participants (and in some cases, over 200 participants). Each participant is assigned a specific role, usually as part of one or more teams. The relationships between teams can be complex and fluid, as they may be rivals or allies, or even have orthogonal or evolving goals. Participants seek to accomplish their goals by engaging with both well defined mechanisms (often involving resources or conflict) as well as with each other (in negotiation of trade and alliances).

Many megagames also incorporate significant amounts of imperfect information or information asymmetry. While information can be broadcast by the facilitators, it is more often given out piecemeal. Members of a team who witnessed (or caused) an event are responsible for disseminating knowledge from that event to the rest of the team members, and privileged knowledge may be used for leverage or tactical advantage. The chaotic nature of play compounds the difficulty of communication, where each participant has limited time to interact with each other, confer with teammates, and come to a decision about a pressing issue.

Games provide a space in which participants can experiment with actions that may be forbidden in normal society (for example, conflict and betrayal). Huizinga (1971) calls this "the magic circle" within which the rules and reality of the game supersede those of everyday life, and within which the consequences of your actions are judged according to different standards. Koljonen (2015) calls this "an alibi for interaction" where the game creates "a situation that allows someone to try a new behavior or experience or initiate an interaction". This magic circle is key in permitting participants the full range of strategic options in the game.

Megagames have also been used as planning exercises by private and governmental organizations, where they are often called simulations, exercises, or wargames. For example, they have been used to test alternative business practices and emergency preparedness (Grayston 2022) as well as planning for communities with aging populations (McCall et al. 2020). Wargames are also used widely for the study of historical and current conflict (Brynen and Milante 2013) and military planning[2].

[1]Megagames are notoriously difficult to define and is an ever evolving genre. South West Megagames, one of the most popular megagame creators, has their definition here: https://www.swmegagames.co.uk/what-is-a-megagame.

[2]See the Connections series of conferences devoted to wargaming in practice: https://connections-wargaming.com

Motivation

While megagames are a direct metaphor for decentralized multiagent systems, our primary motivation for using megagames for pedagogy stems from the benefits of experiential learning highlighted above. Megagames are fun and memorable, and memorable experiences foster durable learning. Moreover, while we teach simple games like Prisoner's Dilemma in the classroom, and make use of simple demos, their actions within these games are generally not driven by real stakes. By embedding the learning experience within a megagame, the fiction of the megagame provides learners with meaningful stakes that are consequential within the magic circle, where participants can care passionately about the outcomes of their decisions. Thus, we leverage the intrinsic motivation within megagames to bring the pedagogical experience one step closer to the real world. Finally, the chaos of play also reminds learners of the limitations of game theory. Interactions with human actors and the messiness of the real world can complicate an otherwise straightforward analysis.

Methodology

The megagame is a full-day activity that takes place after the completion of the lecture component of the course. We use the commercially available product *Den of Wolves* (Mizon 2019), designed to support 40-43 participants, with the help up 3-5 facilitators (the instructor and volunteers). This product was chosen for ease of preparation and facilitation based on reviews and consultation with veteran facilitators[3]. Its core mechanics of resource allocation (summarized below) additionally supports concepts from the course.

Den of Wolves

In *Den of Wolves*, participants assume the role of leaders of a ragtag group of spaceships led by an aging battle cruiser, on the run from an unknown group of enemies called the *Wolves* (managed by the facilitators)[4]. For brevity[5], we describe only the central mechanism for the game as it is the most relevant one to our pedagogical goal.

Each ship must manage its own needs and limited resources (such as food and water) and facilities for resource production. This is further complicated by staffing requirements to operate those facilities (see Fig 1), as well as repairs and recovery necessary in the wake of attacks. For example, the hydrogen mining vessel has facilities for producing water and a dedicated transport shuttle for moving resources, but limited ability to grow food and no provisions for repairs, recovery, or self-defense. The ships are set up so that trade and cooperation are necessary for the smooth running of ships, but the leaders of each ship are in implicit conflict over the very limited resources available within the fleet.

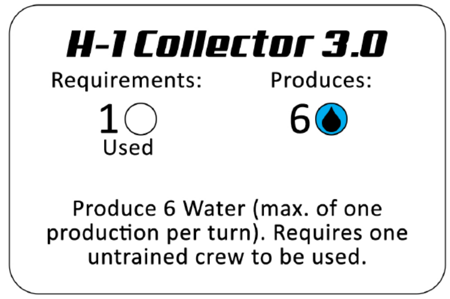

Figure 1: A facility from *Den of Wolves* which requires 1 untrained crew to operate, generating 6 water resources.

This mechanism, taken in aggregate across the fleet, forms a resource allocation optimization problem. But the setup of the game and the chaotic nature of play obscures this from the participants. Moreover, facilitators can (and should) throw additional wrinkles to the players: the fleet must defend against periodic attacks from the Wolves, and confront randomized crisis events. Additionally, the ships are divided between fictitious nations, which have representation within a Model UN style minigame.

Preparation

Our facilitation team consisted of the instructor and three volunteers, one of which was a local expert on megagames with ample experience in their facilitation. Training, as a result, consisted primarily of reading through the PDF files and two video calls, mainly for event logistics.

The megagame product contains many PDF files, both instructions for facilitators and briefings to the participants (individualized for each ship). Moreover, there are PDF files for the variety of mechanical and thematic elements that must be printed out and sorted ahead of time[6].

The megagame event is first announced to students with a brief message in the syllabus and at the start of the course about an optional experiential learning activity, held on the weekend. As it took place out of class, we ensure students are aware of the opportunity as early as possible. This is followed by a second announcement about 3 weeks before the event once event space is finalized and a registration system set up. Because this megagame calls for a fairly strict 40-43 participant count, and we had about 20 students potentially interested, we also open the event up to the general public to fill the remaining spaces (including a small list of alternates). Public tickets were priced according to local community expectations, so that the event itself was revenue neutral. Students were given a registration code that waived the event fee. The inclusion of members of the public is an important element of the activity as it compounds the chaos of play and emphasizes the decentralized nature of the experience.

[3]A list of megagame products is maintained at https://www.megagameassembly.com/ready-made-megagames, and a repository of free megagames at https://www.megagameassembly.com/free-megagames

[4]Science fiction connoisseurs may recognize the plot from *Battlestar Galactica*, an inspiration for the product.

[5]And also to maintain the integrity of the first-time play experience for the adventurous reader.

[6]While we had printouts from a previous run of the game, one can expect to spend 2 afternoons on this task alone, plus additional time to upgrade the materials (with foam tokens, 3D printed figurines, lanyards and name tags, and laminated printouts) as desired.

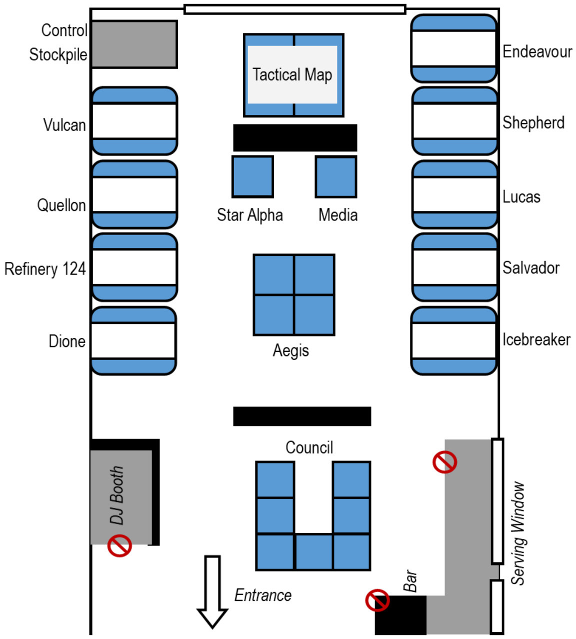

Figure 2: Layout of the event area. Participant zones in blue, with the names of ships or activities. A no-entry sign marks off-limits areas.

Participants may request to be grouped together, which we do where possible. Remaining participants were placed into teams, preferentially grouping students with each other. Teams were then assigned to ships, with highly specialized roles preferentially given to members of the public[7], and then sent their briefing materials which include a high level description of their ship and goals.

Accessibility and Safety

For maximum accessibility, the event space is chosen to be a wheelchair accessible area of the university campus. Additionally, an alternative independent study assignment is available to students who could not or choose not to participate in the event.

During the event, one of the facilitators acts as a safety officer and is able to step out of the "magic circle" to address any players who appear stressed or in distress. For larger events, it is recommended that a facilitator take on the role of safety officer as a primary duty.

Running the Game

Facilitators arrive before the start time to prepare the location. This primarily involves arranging tables and chairs to facilitate the appropriate activities (see Fig 2), and sorting the printouts, cards, and other materials.

Once participants arrive, play begins after a brief welcome introducing the facilitators. Play in megagames typically alternate between an Action Phase where participants move around the play area and negotiate and trade, and a Team Phase where participants return to their teams to discuss next steps. Because of the chaos of play, it is often the case that some members of the team are not aware of events happening elsewhere in the game until being briefed in the Team Phase. This may also be the case for the facilitators as well, and the Team Phase gives the facilitators time to discuss plans (such as Wolf attacks and upcoming crisis events).

We follow best practice in the megagame community in concluding the game. To avoid a mad scramble in the last turn of the game[8], the end of the game is called at the start of a round without prior indicators. We follow the announcement with a round of appreciation for the volunteers and participants, and also there is an explicit reminder that we are stepping now out of the high-tension roles that we have adopted during play (i.e. out of the magic circle), and back into our everyday selves. We then have an informal debriefing session with snacks where participants and facilitators can discuss what happened during the game. Due to the chaos of play, participants, and even facilitators, often only see part of a chain of events, and this debriefing is necessary to appreciate the full chain of causality.

Activity Assessment

An assignment accompanies participation in the activity, with one submission expected from every ship-team with students. The assignment is summative in nature in that it allows students to demonstrate their command of course materials. It contains open-ended questions that asks the students to reflect on their decisions during the game and express them using the language of game theory and multiagent systems (see Supplementary Materials for an example assignment). Each student is asked to contribute one paragraph describing how they made decisions during one or two particularly important events during the game. And the group is asked to reflect on what they would do differently *ex post* – i.e. knowing other teams' preferences or other information unrevealed at the time.

Students were graded in based on the accuracy and clarity of their descriptions, as well as effective use of game theory and multiagent systems terminology. Importantly, students were not graded based on their performance in the megagame, but rather on their ability to articulate their strategy and actions in achieving goals within the game, using the language of the course. Students were additionally invited to submit optional and anonymous surveys through the course portal.

Results

Megagame Outcomes

We ran two events using *Den of Wolves*, corresponding to two cohorts of students of the course in Fall 2022 (23 participants) and 2023 (19 participants). The two events concluded with radically different outcomes. In the first event (see quotations below with the 'D' prefix), miscommunication between the battle cruiser and the other ship leaders led to a

[7]This is to maintain some degree of homogeneity in the situations students would be confronting.

[8]This is a well-known behavior in game theory, as finitely and infinitely repeated games admit radically different solutions.

breakdown of cooperation. Wolf attacks inflicted significant infrastructure damage throughout the fleet. Many ships felt the battle cruiser was coercing the other ships for protection fees, which led to a schism emerging within the fleet, finally resulting in the main resource production vessels splitting off into their own faction. As is best practice in megagame hosting, we concluded the game at this point, leaving the fates of the splinter fleets unresolved[9].

In the second event (quotations with the 'E' prefix), the participants eventually identified the nature of the resource allocation problem and "socialized" the production of critical resources via a centralized mechanism managed by the battle cruiser. While there were some complaints, there was no actual dissent and after a few more rounds, it was clear that the socialized mechanism was effective and we wrapped the game up with congratulations to the fleet.

Survey Results

Because the surveys were optional and hosted online (thus distributed separately from the event), we received few survey responses (7 across both events). This turnout was further confounded by the distribution of a paper survey after the activity that focused on improvement of the megagame itself[10]. While the survey results showed generally positive responses to pedagogical outcomes (see Table 1), we apply qualitative research techniques to discover what pedagogical goals were accomplished by the activity.

Qualitative Analysis

We apply reflexive thematic analysis (RTA) (Braun and Clarke 2021) to the submitted student assignments with the goal of understanding how students used the tools of multiagent system and game theory to make sense of the megagame. RTA is a tool widely used by human-computer interaction (HCI) researchers and social scientists to study meaning-making in various human activities. RTA is a 6-phase iterative process where the data items (for example, transcripts of interviews; or in our case, text of submitted assignments) are first distilled into codes that capture the essential meaning of the passages. Then, from the codes, we generate themes that capture common patterns of meaning and thought repeated across data items. These themes are then translated into the final report.

Reflexive thematic analysis, in particular, embraces the constructivist idea that the researcher's perspective is intrinsically tied to and reflected in the knowledge that is produced. This makes it an appropriate tool for this analysis since, as the instructor and a facilitators at the event, the researcher is inextricably a part of the meaning-making process of the students.

We begin by defining the following research question:

> **RQ1**: How did the students make use of game theory and multiagent systems terminology to make meaning of decisions in the megagame activity?

[9] In the survey comments, several students expressed the desire to play out the splinter fleets.

[10] This is standard practice within the megagame community.

We use a deductive coding technique, where we use existing theory (the language of multiagent systems and game theory) as a starting point. We code at a surface, semantic level, focusing on the literal expressions in the write-ups and what they say about students' experiences. This approach examines whether the student were able to express their thoughts using the language of course, rather than speculating what their thought processes were, which would risk projecting an understanding of the course material where other thought processes were used. We treat each submitted assignment (which may have $1 - 4$ student authors) as one data item. We take one major departure from the RTA process: Where RTA advocate against themes that are simply "topic summaries", we do specifically include several, as the objective of the qualitative analysis is to identify evidence of successful integration of concepts from the course.

Pedagogical Effects

In this section, extracts from the thematic analysis showcase how students successfully use game theory and multiagent system to understand the megagame. All bold text is added to highlight the terms in action.

Effective Use of Game Theory Students used a variety of game theory concepts to guide their decision-making during the megagame. One popular analogy is the Prisoner's Dilemma and its repeated variant as a metaphor for long term cooperation:

> "We decided to follow the **TFT strategy** simply because at first, we thought everyone is truthful with their demands, requirements, and resources. [...] (However,) TFT only seems to work well if we have information about everyone's **incentives**." (D2)

Understanding the fragility of equilibria in repeated Prisoner's Dilemma, one group made use of formal contracts to deter defection:

> "I also realized that (their) **dominant strategy** would be to **defect** and not give us 1 water per turn after we [...] gave them the research, so we made them sign a contract, legally bounded and enforced by the (battle cruiser) and the Council [...]" (D4)

Alternatively, another group reports engaging in the ultimatum game after another ship defaulted on an agreement:

> "Interrupting this core council meeting, I proclaimed our **ultimatum**: No fuel for anyone until we were provided (with our owed resources)." (E7)

Several ships detailed auction protocols that they devised to sell off resources they produced to the highest bidder. In the case of (D5), their ship had no naive resource production, and instead acted as a market-maker for various resources:

> "I told representatives of each ship that if they gave some other game units (e.g. food, water, or fuel), I could give them an untrained crew unit in the next turn [...] This process is an example of an **incentive-compatible auction** where the payments were resources the players were willing to exchange for the untrained crew." (D5)

Question	Likert Scale 1	2	3	4	5
Q1. I felt the experience helped me contextualize some of the concepts from the course.	0	1	0	3	3
Q2. I appreciated the presence of members of the public during the activity.	1	1	0	1	4

Table 1: Responses to pedagogically relevant questions from class survey (Scale: 1 = strongly disagree, 5 = strongly agree)

In some cases, the performance of the auction was improved by strategically hording resources:

> "(By) stockpiling some of our resources, we created a shortage of water in a critical round after Wolves attacked (damaging water production). I (inferred) other teams would highly value the water we produced during the shortage [...] By holding strong in resource trade negotiations, we were able to capitalize on the now high value of water for even more resources" (E5)

Finally, students in the Council were able to engage in strategic voting and strategic timing of votes:

> "We approached the (nation A's) representative [...] and suggested that we pool our votes to elect a citizen of (nation B) who would provide positive coverage of our nations. [...] we would win at least a **plurality** of the votes" (E3)

> "I waited for other representatives to cast their votes so I would not go against a specific alliance, especially when my vote would not be the **deciding vote**." (D1)

References to Game Theory Terminology Many submissions also used game theory terminology to make sense of the situations the students found themselves in:

> "we tried our best (we **cooperated** with (the Council)) even though some of their decisions generated **envy** on our part as they did not take us into account when making **allocations of resources**" (D2)

> "(We) pointed out [...] that they were "**defecting**" on the fleet by not thinking about the Greater Good for the entire fleet (**Utilitarian** strategy), since staying in the battle was causing damage. Instead, they were selfishly threatening **Mutually Assured Destruction** (MAD) should they not get paid" (D1)

Coercion and Monopoly The "magic circle" of the game allowed more extreme strategies to emerge. In the first game, the students on board the battle cruiser experimented with trading protection during Wolf attack for resources. This was received poorly by the rest of the fleet:

> "At the time, it seemed like a **Pareto optimal** situation. I predicted that other nations lacking fighter squadrons and any proper defenses would think of it in the same way as I did. But other nations turned hostile to this request of mine and considered it blackmail[11]" (D2)

[11] Students from several group referred to this coercive behavior incorrectly as "blackmail"; we assume this accusation was leveled during the game and it stuck.

Interestingly, the ship that had a monopoly on fuel production successfully leveraged their monopoly by timing the market. With an appropriate spin, their service was perceived as "responsible":

> "we realized that the more fuel other ships had, the less they also valued fuel and so it became harder to make trades for vital resources [...] It was at that point, we decided to severely limit the amount of fuel in circulation [...] This gave us two key advantages, one was that the other ships liked us since we were responsible for keeping them alive with jumps, and two was that we were able to keep fuel valuable throughout the game" (E7)

Preference for Centralized Mechanisms Perhaps informed by their lessons on fair division, several students in both games express a desire for centralizing the mechanism for allocating resources. The battle cruiser in the first game laid the ground work for this early on:

> "Then, to develop good strategies we needed both to know what we needed and what the other teams needed (other team's **valuations** and **utilities**). [...] As a next step, I went to each of the other teams to make a plan of the fleet: where was each vessel and what did they have that would be of interest to us (sort of utility map)." (D2)

Unfortunately, perceptions of coercion ousted them from a position of trust:

> "At some point, there was a discussion to **centralize** resource control. This was before everyone completely lost their trust in the (battle cruiser)" (D1)

Other ships also envisioned similar ideas (naturally with themselves as the trusted mechanism):

> "I believe we could have remedied this issue (selfish and dishonest demands for resources) by coming up with a **mechanism** such that all ships would be able work together in an orderly manner and would receive the most utility. This would have involved having each ship send someone to (our ship) to state their needs and capabilities, and having (our ship) process all this information and decide on what each ship should do." (D5)

The central mechanism was successfully implemented in the second game for some resources. Naturally, students of ships whose resource production were "socialized" chaffed at their diminished economic leverage:

> "At council I did not want to disagree with this socialization too much because I would lose reputation

> if I did, however I did want to argue against it somewhat as by giving away control of our main resource we would lose the **utility** we could gain by bartering with that resource" (E1)

Chaos of Play One pedagogical goal was to demonstrate the challenges one might face implementing more sophisticated techniques from class in the chaotic environment of play, limited by time and with imperfect information. The former is evidenced by the following quotations:

> "Had there been more time [...], it could have been interesting to **collude** with the other shuttle users to form a shuttle service corporation, where we only rent out a limited number of shuttles each turn (like De Beers diamonds), and sell them off in a second price sealed bid type auction" (D3)

> "The reason why **English Auction** is not used is that our time is limited and we can only conduct one round of auction." (D4)

There were also numerous opportunities for information to be lost or delayed in the chaos:

> "Due to the chaotic nature of the game, people were not able to fully grasp the situation of the entire fleet, and so acted on the **limited information** that they had. This resulted in people acting selfishly and not being **truthful**." (D5)

> "we were not aware of the massive demand for materials across the system, so we continued to produce ore instead." (D1)

> "The bottleneck on training engineers led me to consider **auctioning** the capacity [...], but this was never transmitted over the news and caused concern within the council, so it did not happen." (E6)

Indeed, one student even realized that they had become a tool of active disinformation:

> "I became almost a spokesperson that directed anonymous feedback to the fleet [...] Later the other ships saw the value in anonymous feedback and began to take advantage of this to spread false information and trick other players into gaining additional resources for themselves." (D9)

Finally, one student summed up the experience succinctly using a concept and example from class:

> "This uncertainty [...] meant we were playing a **Basyesian game**. Our interactions were thus similar to the **Sheriff's dilemma** example that we had discussed in class, but with the added complication created by repeated interaction." (E3)

Discussion and Reflection

The turnout for the megagame is very high for an optional full-day activity on a weekend: with 42 joining the activity out of a total of 58 students across both classes.

The previous section showcased examples successful usage of multiagent system or game theory concepts during the activity or in analysis of the aftermath. In addressing RQ1, student descriptions show a surprising variety in the concepts utilized, and in some cases, with surprising depth. Perhaps unsurprisingly, sometimes concepts were also used incorrectly, the terms used is not the one most suitable, or the connection to the concept was tenuous at best.

Chaos of Play An important goal of this activity is to bring the students outside of the sterile theory of class and immerse them in a semblance of the chaotic real world where they can try to put the theory to use. Interacting with other people is rarely as straightforward as playing Cooperate or Defect in a Prisoner's Dilemma, and the lack of information and time is often a constraint. Students' understanding and appreciation for these constraints are well demonstrated above in the results of the thematic analysis. The inclusion of members of the public also added to this unpredictability. While, students clearly recognized different behavior in the activity from members of the public (data not shown), the few survey results (see Table 1) – and the lack of anecdotal complaints – showed that they were a positive addition.

Long-term Benefits We also observed unexpected long-term effects across the years. Since the megagame is scheduled at around the same time each year, it became a highly anticipated event within the university community. Rumors from the first game would circulate to second cohort (though we made some changes to subvert these expectations, and in one case, the 'tale' from the first game cited in a write-up was demonstratively incorrect). However, the discovery of a centralized solution in the second game required a different game had to be selected for the following year. The popularity of the game also led to improved interest in the course (evidence is anecdotal as the course was already at full capacity), and in a few cases, this interest translated into students taking on extracurricular research project or staying on for graduate school.

Conclusion and Future Work

We conducted experiential learning in a large activity in a AI course to help students learn theoretical concepts through experience. We show, using an assignment with open-ended questions, that we can elicit effective use of course concepts. Moreover, the "magic circle" of the game environment allows participants to experience game theory in action, with meaningful, fictitious stakes. The activity is very well received by students and has led to enduring interest in the domain and long-term research partnerships.

Megagames are highly customisable and offer a flexible medium with which to provide experiential learning to students. While we have chosen a science fictional domain, megagames can be selected or created to feature domains of pedagogical interest in computer science – for example, AI safety and ethics, sustainable computation, privacy and security, and government regulation of AI development. Such an activity can accomplish the dual goal of facilitating education in both game theory education and in the featured domain.

Acknowledgments

This activity would not be possible without the generous efforts of our volunteers from Ottawa Megagames, especially Kevin Farnworth, Forest Anderson, and Lianne Davies. This study was approved by the Carleton University Research Ethics Board-B Project #124657.

References

Abraham, D. J.; Blum, A.; and Sandholm, T. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*, 295–304.

Baltes, J.; Sklar, E.; and Anderson, J. 2004. Teaching with RoboCup. *Accessible Hands-on Artificial Intelligence and Robotics Education*, 131–135.

Braun, V.; and Clarke, V. 2021. *Thematic analysis: A practical guide*. SAGE Publications Ltd.

Brynen, R.; and Milante, G. 2013. Peacebuilding with games and simulations. *Simulation & Gaming*, 44(1): 27–35.

Burguillo, J. C. 2010. Using game theory and competition-based learning to stimulate student motivation and performance. *Computers & Education*, 55(2): 566–575.

Deci, E. L.; and Ryan, R. M. 2013. *Intrinsic motivation and self-determination in human behavior*. Springer Science & Business Media.

Ebner, M.; and Holzinger, A. 2007. Successful implementation of user-centered game based learning in higher education: An example from civil engineering. *Computers & education*, 49(3): 873–890.

Geleta, M.; Xu, J.; Loya, M.; Wang, J.; Singh, S.; Li, Z.; and Gago-Masague, S. 2023. Maestro: A gamified platform for teaching AI robustness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 15816–15824.

Grayston, R. 2022. Megagames at Work - Rob Grayston - MegaCon 2022 Seminars. https://www.youtube.com/watch?v=YYhRerL84Oc. Accessed: 2025-08-01.

Huizinga, J. 1971. *Homo ludens: A study of the play-element in culture*. Beacon press.

Kahng, A.; Mackenzie, S.; and Procaccia, A. 2021. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research*, 70: 1223–1252.

Kolb, D. A. 2014. *Experiential learning: Experience as the source of learning and development*. FT press.

Koljonen, J. 2015. Introduction to Narrative Experience Design. Opening Lecture at Alibis for Interaction, Landskrona, Sweden. https://www.youtube.com/watch?v=yXjEHDBjrXE.

Lee, H.-I.; Ignatyev, D.; Shin, H.-S.; and Tsourdos, A. 2024. Project-Based Learning Course Design for Multi-Agent Autonomy Using Quadrotors. *IFAC-PapersOnLine*, 58(16): 107–112.

Lee, J. J.; and Hammer, J. 2011. Gamification in education: What, how, why bother? *Academic exchange quarterly*, 15(2): 146.

Lepper, M. R. 1988. Motivational considerations in the study of instruction. *Cognition and instruction*, 5(4): 289–309.

McCall, V.; Ziegler, F.; Robertson, J.; Lovatt, M.; Phillips, J.; Porteus, J.; McIntyre, Z.; Rutherford, A.; Sixsmith, J.; Woolrych, R.; et al. 2020. Housing and ageing: Let's get serious—"How do you plan for the future while addressing immediate chaos?". *Social Inclusion*, 8(3): 28–42.

Meir, R. 2018. *Strategic voting*. Morgan & Claypool Publishers.

Mizon, J. 2019. Den of Wolves. https://www.swmegagames.co.uk/den-of-wolves.

Reigeluth, C. M. 2013. *Instructional-design theories and models: A new paradigm of instructional theory, Volume II*. Routledge.

Rosenschein, J. 2025. Multiagent Systems and the Search for Appropriate Foundations. Keynote at the 24th International Conference on Autonomous Agents and Multiagent Systems, Detroit, USA.

Scarlatos, L. L.; Tomkiewicz, M.; and Courtney, R. 2013. Using an Agent-Based Modeling Simulation and Game to Teach Socio-Scientific Topics. *IxD&A*, 19: 77–90.

Tohidi, H.; and Jabbari, M. M. 2012. The effects of motivation in education. *Procedia-social and behavioral Sciences*, 31: 820–824.

Tsang, A.; and Larson, K. 2017. If you like it, then you shoulda put a sticker on it. In *EXPLORE Workshop, International Conference on Autonomous Agents and Multi-Agent Systems*.

United Nations. n.d. Model United Nations — United Nations. https://www.un.org/en/mun.

Vidal, J. M.; and Buhler, P. 2002. Using RoboCup to teach multiagent systems and the distributed mindset. In *Proceedings of the 33rd SIGCSE technical symposium on Computer science education*, 3–7.

Ye, J.; Li, Y.; Zou, W.; and Wang, X. 2025. From Awareness to Action: The Effects of Experiential Learning on Educating Users about Dark Patterns. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–22.